

# **Exhibit 9**

**Exhibit 9 to Complaint**  
**Intellectual Ventures I LLC and Intellectual Ventures II LLC**


**Example Comerica Count IV Systems and Services**  
**U.S. Patent No. 7,949,785 (“the ’785 Patent”)**

The Accused Systems and Services include, without limitation, Comerica’s systems that utilize Apache Spark (“Spark”); all past, current and future systems and services that operate in the same or substantially similar manner as the specifically identified systems and services; and all past, current and future Comerica’s systems and services that have the same or substantially similar features as the specifically identified systems and services (“Example Comerica Count IV Systems and Services”).

<b>U.S. Patent No. 7,949,785</b>	
<b>Example Claim 30</b>	<b>Accused Instrumentalities: Example Comerica Count IV Systems and Services</b>
30. A virtual network manager, comprising:	<p>Upon information and belief, the Comerica's systems include "a virtual network manager." The virtual network manager is the Kubernetes functionality related to DNS for Services and Pods.</p> <p><b><u>DNS for Services and Pods</u></b></p> <p><u>Kubernetes creates DNS records for Services and Pods.</u> You can contact Services with consistent DNS names instead of IP addresses.</p> <p>Kubernetes publishes information about Pods and Services which is used to program DNS. Kubelet configures Pods' DNS so that running containers can lookup Services by name rather than IP.</p> <p><u>Services defined in the cluster are assigned DNS names.</u> By default, a client Pod's DNS search list includes the Pod's own namespace and the cluster's default domain.</p> <p><i>See <a href="https://kubernetes.io/docs/concepts/services-networking/dns-pod-service/#pod-s-dns-config">https://kubernetes.io/docs/concepts/services-networking/dns-pod-service/#pod-s-dns-config</a></i> (last accessed on November 11, 2023).</p> <p><u>A Service in Kubernetes is an abstraction which defines a logical set of Pods</u> and a policy by which to access them. Services enable a loose coupling between dependent Pods. A Service is defined using YAML or JSON, like all Kubernetes object manifests. The set of Pods targeted by a Service is usually determined by a <i>label selector</i> (see below for why you might want a Service without including a <i>selector</i> in the spec).</p> <p><i>See <a href="https://kubernetes.io/docs/tutorials/kubernetes-basics/expose/expose-intro/">https://kubernetes.io/docs/tutorials/kubernetes-basics/expose/expose-intro/</a></i> (last accessed on November 11, 2023).</p>

	<p>Kubernetes assigns this Service an IP address (the <i>cluster IP</i>), that is used by the <u>virtual IP address mechanism</u>. For more details on that mechanism, read <a href="#">Virtual IPs and Service Proxies</a>.</p> <p>See <a href="https://kubernetes.io/docs/concepts/services-networking/service/">https://kubernetes.io/docs/concepts/services-networking/service/</a> (last accessed on November 11, 2023).</p> <h2><b><u>Virtual IPs and Service Proxies</u></b></h2> <p>Every node in a Kubernetes cluster runs a <code>kube-proxy</code> (unless you have deployed your own alternative component in place of <code>kube-proxy</code>).</p> <p>The <code>kube-proxy</code> component is responsible for implementing a <i>virtual IP</i> mechanism for Services of type <code>other than ExternalName</code>.</p> <p>See <a href="https://kubernetes.io/docs/reference/networking/virtual-ips/">https://kubernetes.io/docs/reference/networking/virtual-ips/</a> (last accessed on November 11, 2023).</p> <p>Comerica's systems include Kubernetes functionality. <i>See, e.g.</i>, <a href="https://careers.comerica.com/job/19419241/principal-engineer-cloud-automation-governance-and-engineering-cage-dallas-or-detroit-metro-auburn-hills-mi/">https://careers.comerica.com/job/19419241/principal-engineer-cloud-automation-governance-and-engineering-cage-dallas-or-detroit-metro-auburn-hills-mi/</a> (last accessed on November 13, 2023).</p>
a network interface configured for data communication via a virtual network that is defined by a domain name having an associated public network address;	Upon information and belief, the Comerica's systems utilize Kubernetes and include “a network interface configured for data communication via a virtual network that is defined by a domain name having an associated public network address.” For example, Kubernetes DNS for Services and PODS is implemented in part through a network interface called <code>kube-proxy</code> , which maintains network rules that allow communication to pods from network sessions. <code>Kube-proxy</code> communications can relate to network sessions inside or outside of the cluster (the virtual network) in which the services are defined. The services defined in the cluster are assigned DNS names having an associated public network address.

## Kubernetes Components



### kube-proxy


kube-proxy is a network proxy that runs on each node in your cluster, implementing part of the Kubernetes Service concept.

kube-proxy maintains network rules on nodes. These network rules allow network communication to your Pods from network sessions inside or outside of your cluster.

See <https://kubernetes.io/docs/concepts/overview/components/> (last accessed on November 11, 2023).

	<p><b>DNS</b></p> <p>A cluster-aware DNS server, such as CoreDNS, watches the Kubernetes API for new Services and creates a set of DNS records for each one. If DNS has been enabled throughout your cluster then <u>all Pods should automatically be able to resolve Services by their DNS name.</u></p> <p><b>External IPs</b></p> <p>If <u>there are external IPs that route to one or more cluster nodes, Kubernetes Services can be exposed on those externalIPs</u>. When network traffic arrives into the cluster, with the external IP (as destination IP) and the port matching that Service, rules and routes that Kubernetes has configured ensure that the traffic is routed to one of the endpoints for that Service.</p> <p><i>See <a href="https://kubernetes.io/docs/concepts/services-networking/service/">https://kubernetes.io/docs/concepts/services-networking/service/</a> (last accessed on November 11, 2023).</i></p> <p><b>DNS for Services and Pods</b></p> <p>Kubernetes creates DNS records for Services and Pods. You can contact <u>Services with consistent DNS names instead of IP addresses</u>.</p> <p><u>Services defined in the cluster are assigned DNS names.</u> By default, a client Pod's DNS search list includes the Pod's own namespace and the cluster's default domain.</p> <p><i>See <a href="https://kubernetes.io/docs/concepts/services-networking/dns-pod-service/#pod-s-dns-config">https://kubernetes.io/docs/concepts/services-networking/dns-pod-service/#pod-s-dns-config</a> (last accessed on November 11, 2023).</i></p>
--	--

<p>a memory and a processor to implement a register module configured to register devices in a virtual network</p>	<p>Upon information and belief, the Comerica's systems utilize Kubernetes and includes "a memory and a processor to implement a register module configured to register devices in a virtual network." A service runs on a set of pods (devices). The Kubernetes DNS Service module (register module) watches the Kubernetes API for incoming new services and creates a set of DNS records for each of the pods associated with the service (register devices).</p> <p><i>Kubernetes runs your workload by placing containers into Pods to run on Nodes. A node may be a virtual or physical machine, depending on the cluster. Each node is managed by the control plane and contains the services necessary to run Pods.</i></p> <p><i>See <a href="https://kubernetes.io/docs/concepts/architecture/nodes">https://kubernetes.io/docs/concepts/architecture/nodes</a> (last accessed on November 11, 2023).</i></p> <p><b>Initializing your control-plane node</b></p> <p><i>The control-plane node is the machine where the control plane components run, including etcd (the cluster database) and the API Server (which the kubectl command line tool communicates with).</i></p> <p><i>See <a href="https://kubernetes.io/docs/setup/production-environment/tools/kubeadm/create-cluster-kubeadm/">https://kubernetes.io/docs/setup/production-environment/tools/kubeadm/create-cluster-kubeadm/</a> (last accessed on November 11, 2023).</i></p> <p><b>DNS</b></p> <p><i>You can (and almost always should) set up a DNS service for your Kubernetes cluster using an add-on.</i></p> <p><i>A cluster-aware DNS server, such as CoreDNS, watches the Kubernetes API for new Services and creates a set of DNS records for each one. If DNS has been enabled throughout your cluster then all Pods should automatically be able to resolve Services by their DNS name.</i></p> <p><i>For example, if you have a Service called <code>my-service</code> in a Kubernetes namespace <code>my-ns</code>, the control plane and the DNS Service acting together create a DNS record for <code>my-service.my-ns</code>. Pods in the <code>my-ns</code> namespace should be able to find the service by doing a name lookup for <code>my-service</code> (<code>my-service.my-ns</code> would also work).</i></p> <p><i>See <a href="https://kubernetes.io/docs/concepts/services-networking/service">https://kubernetes.io/docs/concepts/services-networking/service</a> (last accessed on November 11, 2023).</i></p>
--	---

	<h2>Services in Kubernetes</h2> <p>The Service API, part of Kubernetes, is an abstraction to help you expose groups of Pods over a network. Each Service object defines a logical set of endpoints (usually these endpoints are Pods) along with a policy about how to make those pods accessible.</p> <p>For example, consider a stateless image-processing backend which is running with 3 replicas. Those replicas are fungible—frontends do not care which backend they use. While the actual Pods that compose the backend set may change, the frontend clients should not need to be aware of that, nor should they need to keep track of the set of backends themselves.</p> <p>The Service abstraction enables this decoupling.</p> <p>The set of Pods targeted by a Service is usually determined by a selector that you define. To learn about other ways to define Service endpoints, see <a href="#">Services without selectors</a>.</p> <p>See <a href="https://kubernetes.io/docs/reference/kubectl/cheatsheet/">https://kubernetes.io/docs/reference/kubectl/cheatsheet/</a> (last accessed on November 11, 2023).</p>
the register module further configured to: receive a registration request from an agent associated with a device;	<p>Upon information and belief, the Comerica's systems include “the register module further configured to receive a registration request rom an agent associated with a device.” A service runs on a set of pods. The Kubernetes DNS Service module (register module) watches the Kubernetes API for incoming new services (registration requests) being made via tools like kubectl (agent associated with a device).</p> <div data-bbox="946 833 1586 1034">  <pre> Create a new ClusterIP service named my-cs kubectl create service clusterip my-cs --tcp=5678:8080  Create a new ClusterIP service named my-cs (in headless mode)  kubectl create service clusterip my-cs --clusterip="None" </pre> </div> <p>See <a href="https://kubernetes.io/docs/setup/production-environment/tools/kubeadm/create-cluster-kubeadm/">https://kubernetes.io/docs/setup/production-environment/tools/kubeadm/create-cluster-kubeadm/</a> (last accessed on November 11, 2023).</p>

## Discovering services

For clients running inside your cluster, Kubernetes supports two primary modes of finding a Service: environment variables and DNS.

### DNS

You can (and almost always should) set up a DNS service for your Kubernetes cluster using an add-on.

A cluster-aware DNS server, such as CoreDNS, watches the Kubernetes API for new Services and creates a set of DNS records for each one. If DNS has been enabled throughout your cluster then all Pods should automatically be able to resolve Services by their DNS name.

For example, if you have a Service called `my-service` in a Kubernetes namespace `my-ns`, the control plane and the DNS Service acting together create a DNS record for `my-service.my-ns`. Pods in the `my-ns` namespace should be able to find the service by doing a name lookup for `my-service` (`my-service.my-ns` would also work).

See <https://kubernetes.io/docs/concepts/services-networking/service> (last accessed on November 11, 2023).

## Services in Kubernetes

The Service API, part of Kubernetes, is an abstraction to help you expose groups of Pods over a network. Each Service object defines a logical set of endpoints (usually these endpoints are Pods) along with a policy about how to make those pods accessible.

## Defining a Service

A Service is an object (the same way that a Pod or a ConfigMap is an object). You can create, view or modify Service definitions using the Kubernetes API. Usually you use a tool such as `kubectl` to make those API calls for you.

See <https://kubernetes.io/docs/concepts/services-networking/service> (last accessed on November 11, 2023).

	<pre>kubectl logs deploy/my-deployment kubectl logs deploy/my-deployment -c my-container</pre> <p># dump Pod Logs for a Deployment (single-container case) # dump Pod Logs for a Deployment (multi-container case)</p> <p>See <a href="https://kubernetes.io/docs/reference/kubectl/cheatsheet/">https://kubernetes.io/docs/reference/kubectl/cheatsheet/</a> (last accessed on November 11, 2023).</p>
<p>distribute a virtual network address to the device when the device is registered in the virtual network, the device being identified to other devices in the virtual network by the virtual network address; and</p>	<p>Upon information and belief, the Comerica's systems include a register module further configured to "distribute a virtual network address to the device when the device is registered in the virtual network, the device being identified to other devices in the virtual network by the virtual network address." A service runs on a set of pods. A new service is assigned to a cluster IP ("virtual network address"). As part of configuration, the Kubernetes DNS Service module distributes the network address to at least one pod (device) on which the service runs.</p> <p><b>Using a Service to Expose Your App</b></p> <p>Although each Pod has a unique IP address, those IPs are not exposed outside the cluster without a Service. Services allow your applications to receive traffic. Services can be exposed in different ways by specifying a <code>type</code> in the <code>spec</code> of the Service:</p> <ul style="list-style-type: none"> <li>• <code>ClusterIP</code> (default) - Exposes the Service on an internal IP in the cluster. This type makes the Service only reachable from within the cluster.</li> </ul> <p>See <a href="https://kubernetes.io/docs/tutorials/kubernetes-basics/expose/expose-intro/">https://kubernetes.io/docs/tutorials/kubernetes-basics/expose/expose-intro/</a> (last accessed on November 11, 2023).</p> <p><b>Service ClusterIP allocation</b></p> <p>In Kubernetes, <code>Services</code> are an abstract way to expose an application running on a set of Pods. Services can have a cluster-scoped virtual IP address (using a Service of <code>type: ClusterIP</code>). Clients can connect using that virtual IP address, and Kubernetes then load-balances traffic to that Service across the different backing Pods.</p> <p>See <a href="https://kubernetes.io/docs/concepts/services-networking/cluster-ip-allocation/">https://kubernetes.io/docs/concepts/services-networking/cluster-ip-allocation/</a> (last accessed on November 11, 2023).</p>

### VIP

a virtual IP address, such as the one assigned to every Service in Kubernetes

See <https://kubernetes.io/docs/tutorials/services/source-ip/> (last accessed on November 11, 2023).


### DNS

You can (and almost always should) set up a DNS service for your Kubernetes cluster using an [addon](#).

A cluster-aware DNS server, such as CoreDNS, watches the Kubernetes API for new Services and creates a set of DNS records for each one. If DNS has been enabled throughout your cluster then all Pods should automatically be able to resolve Services by their DNS name.

For example, if you have a Service called `my-service` in a Kubernetes namespace `my-ns`, the control plane and the DNS Service acting together create a DNS record for `my-service.my-ns`. Pods in the `my-ns` namespace should be able to find the service by doing a name lookup for `my-service` (`my-service.my-ns` would also work).

See <https://kubernetes.io/docs/concepts/services-networking/service> (last accessed on November 11, 2023).



	<p>See <a href="https://kubernetes.io/docs/tutorials/kubernetes-basics/update/update-intro/">https://kubernetes.io/docs/tutorials/kubernetes-basics/update/update-intro/</a> (last accessed on November 11, 2023).</p>
a DNS server for the virtual network, the DNS server configured to receive a DNS request from a first device in the virtual network and return a network address associated with a network route director,	<p>Upon information and belief, the Comerica's systems include "a DNS server for the virtual network, the DNS server configured to receive a DNS request from a first device in the virtual network." A client/frontend pod ("a first device") sends a DNS query ("DNS Request") using Kubernetes DNS query functionality, such as CoreDNS or kube-dns, which, includes resolver routines associated with the requesting pod. The DNS server (cluster DNS server (e.g. CoreDNS or kube dns)) is configured to return an IP address of cluster DNS server, which is referenced/returned by the process of the kubelet accessing resolv.conf.</p> <pre> resolv.conf(5)           File Formats Manual           resolv.conf(5)  NAME      top          <u>resolv.conf - resolver configuration file</u>  SYNOPSIS  top          <u>/etc/resolv.conf</u>  DESCRIPTION  top          <u>The resolver is a set of routines in the C library that provide</u> <u>access to the Internet Domain Name System (DNS). The resolver</u> <u>configuration file contains information that is read by the</u> <u>resolver routines the first time they are invoked by a process.</u> <u>The file is designed to be human readable and contains a list of</u> <u>keywords with values that provide various types of resolver</u> <u>information. The configuration file is considered a trusted</u> <u>source of DNS information; see the trust-ad option below for</u> <u>details.</u> </pre> <p>See <a href="https://man7.org/linux/man-pages/man5/resolv.conf.5.html">https://man7.org/linux/man-pages/man5/resolv.conf.5.html</a> (last accessed on November 11, 2023).</p>

## Using CoreDNS for Service Discovery

### About CoreDNS


CoreDNS is a flexible, extensible DNS server that can serve as the Kubernetes cluster DNS. Like Kubernetes, the CoreDNS project is hosted by the CNCF.

You can use CoreDNS instead of kube-dns in your cluster by replacing kube-dns in an existing deployment, or by using tools like kubeadm that will deploy and upgrade the cluster for you.

See <https://kubernetes.io/docs/tasks/administer-cluster/coredns/> (last accessed on November 11, 2023).

### Introduction

NodeLocal DNSCache improves Cluster DNS performance by running a DNS caching agent on cluster nodes as a DaemonSet. In today's architecture, Pods in 'ClusterFirst' DNS mode reach out to a kube-dns serviceIP for DNS queries. This is translated to a kube-dns/CoreDNS endpoint via iptables rules added by kube-proxy. With this new architecture, Pods will reach out to the DNS caching agent running on the same node, thereby avoiding iptables DNAT rules and connection tracking. The local caching agent will query kube-dns service for cache misses of cluster hostnames ("cluster.local" suffix by default).



See <https://kubernetes.io/docs/tasks/administer-cluster/nodelocaldns/> (last accessed on November 11, 2023).

#### Pod foo


When a pod sends an API request to a service within the same Kubernetes cluster, it must first resolve the IP address of the service. To do this, the pod performs a DNS lookup using the DNS server specified in its `/etc/resolv.conf` configuration file.

This file, which is provisioned by the Kubelet, defines the settings for DNS lookups in the pod. It contains a reference to the cluster DNS server.

By default, this configuration file looks something like this:

```
search namespace.svc.cluster.local svc.cluster.local cluster.local
nameserver 10.123.0.10
options ndots:5
```

### DNS lookups on services



When a pod performs a DNS lookup, the query is first sent to the local DNS resolver in the pod. This resolver uses the `resolv.conf` configuration file. In this file, the `node.localdns` server is set up as the default recursive DNS resolver, which acts as a cache.

If this cache does not contain the IP address for the requested hostname, the query is forwarded to the cluster DNS server (`CoreDNS`).

This DNS server determines the IP address by consulting the Kubernetes service registry. This registry contains a mapping of service names to their corresponding IP addresses. This allows the cluster DNS server to return the correct IP address to the requesting pod.

Any domains that are queried but are not in the Kubernetes service registry are forwarded to an upstream DNS server.

We will go through each of these components in more detail step-by-step.

See <https://medium.com/@seifeddinerajhi/connecting-the-dots-understanding-how-pods-talk-in-kubernetes-networks-992fa69fbbeb> (last accessed on November 14, 2023).

<p>a private network address associated with a second device in the virtual network, and a virtual network address associated with the second device.</p>	<p>Upon information and belief, the Comerica's systems include "a private network address associated with a second device in the virtual network, and a virtual network address associated with the second device." The private network address associated with a second device is a unique private IP address assigned to a backend pod. The virtual network address associated with the second device is the Cluster IP or Virtual IP assigned to a set of pods.</p>  <p>CoreDNS</p> <p>≡</p> <h2>kubernetes</h2> <p><a href="#">View Source</a></p> <p><i>kubernetes</i> enables reading zone data from a Kubernetes cluster.</p> <h3>Description</h3> <p>This plugin implements the <a href="#">Kubernetes DNS-Based Service Discovery Specification</a>.</p> <p>CoreDNS running the kubernetes plugin can be used as a replacement for kube-dns in a Kubernetes cluster. See the <a href="#">deployment</a> repository for details on <a href="#">how to deploy CoreDNS in Kubernetes</a>.</p>
---	--

	<p> <b>CoreDNS</b></p> <h2>Metadata</h2> <p>The <code>kubernetes</code> plugin will publish the following metadata, if the metadata plugin is also enabled:</p> <ul style="list-style-type: none"><li>• <code>kubernetes/endpoint</code>: the endpoint name in the query</li><li>• <code>kubernetes/find</code>: the resource kind (pod or svc) in the query</li><li>• <code>kubernetes/namespace</code>: the namespace in the query</li><li>• <code>kubernetes/port-name</code>: the port name in an SRV query</li><li>• <code>kubernetes/protocol</code>: the protocol in an SRV query</li><li>• <code>kubernetes/service</code>: the service name in the query</li><li>• <code>kubernetes/client-namespace</code>: the client pod's namespace (see requirements below)</li><li>• <code>kubernetes/client-pod-name</code>: the client pod's name (see requirements below)</li></ul> <p>The <code>kubernetes/client-namespace</code> and <code>kubernetes/client-pod-name</code> metadata work by reconciling the client IP address in the DNS request packet to a known pod IP address. Therefore the following is required:</p> <ul style="list-style-type: none"><li>• <code>pods-verified</code> mode must be enabled</li><li>• the remote IP address in the DNS packet received by CoreDNS must be the IP address of the Pod that sent the request.</li></ul> <hr/> <p>See <a href="https://coredns.io/plugins/kubernetes/">https://coredns.io/plugins/kubernetes/</a> (last accessed on November 11, 2023).</p> <h3>Connecting Applications with Services</h3> <p>Kubernetes assumes that pods can communicate with other pods, regardless of which host they land on. <u>Kubernetes gives every pod its own cluster-private IP address</u>, so you do not need to explicitly create links between pods or map container ports to host ports. This means that containers within a Pod can all reach each other's ports on localhost, and all pods in a cluster can see each other without NAT. The rest of this document elaborates on how you can run reliable services on such a networking model.</p> <p>See <a href="https://kubernetes.io/docs/tutorials/services/connect-applications-service/">https://kubernetes.io/docs/tutorials/services/connect-applications-service/</a> (last accessed on November 11, 2023).</p>
--	---

### Using a Service to Expose Your App

Although each Pod has a unique IP address, those IPs are not exposed outside the cluster without a Service. Services allow your applications to receive traffic. Services can be exposed in different ways by specifying a `type` in the `spec` of the Service:

- ClusterIP (default) - Exposes the Service on an Internal IP in the cluster. This type makes the Service only reachable from within the cluster.

See <https://kubernetes.io/docs/tutorials/kubernetes-basics/expose/expose-intro/> (last accessed on November 11, 2023).

### Service ClusterIP allocation

In Kubernetes, Services are an abstract way to expose an application running on a set of Pods. Services can have a cluster-scoped virtual IP address (using a Service of `type: ClusterIP`). Clients can connect using that virtual IP address, and Kubernetes then load-balances traffic to that Service across the different backing Pods.

See <https://kubernetes.io/docs/concepts/services-networking/cluster-ip-allocation> (last accessed on November 11, 2023).